demonstrates that, in theory, the linkage can be statically balanced for all base link orientations and also provides a three-step link mass-distribution procedure to achieve the balance. Perfect static balance obviates the need for actuator or external forces to support unbalanced linkage weight for any pose, freeing up more force for useful work.

Finally, the three-DOF architecture is scalable from large crane-sized construction and material-handling equipment down to micro-machines for use in applications such as minimally invasive surgery. Arm- and finger-scale three-DOF force-feel joysticks (figure 2) demonstrate a portion of this scalability.

Point of Contact: B. Adelstein
(650) 604-3922
dadelstein@mail.arc.nasa.gov

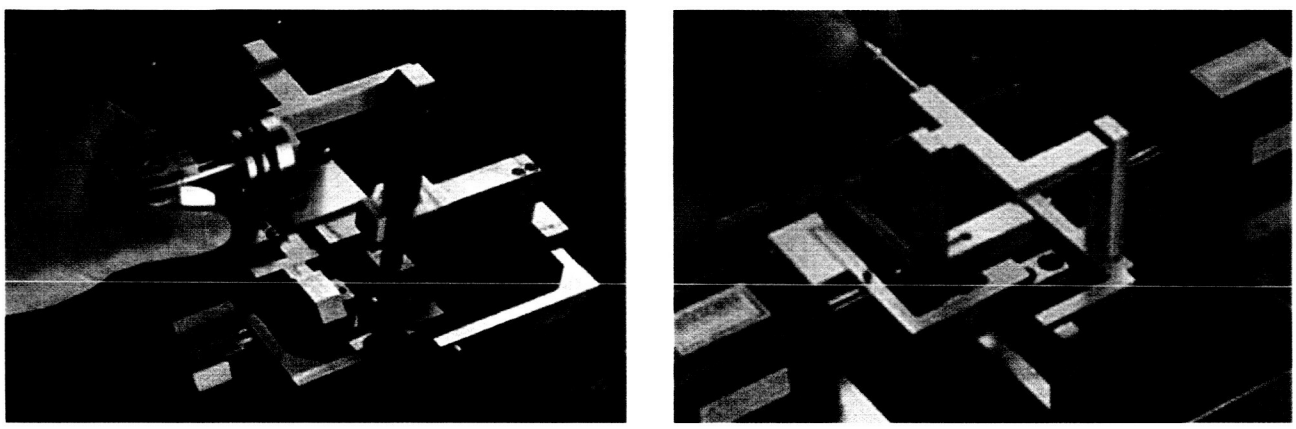

Fig. 2. Arm- and finger-scale force-feel hand-controllers.

Performance of the COSMOS Multi-Level Parallelism Molecular Dynamics Code on the 512 CPU Origin System

James R. Taft

Ames recently purchased an SGI 512 CPU Origin 2000 system. The system has been named Lomax, after the late celebrated Ames researcher Harvard Lomax. The Lomax system is the largest single shared-memory multi-processor system in the world (see figure 1). It is the result of an Ames-driven partnership with SGI to push the limits of single-system shared memory designs. It is believed that large CPU count single-system designs offer many potential advantages in those research areas that require very high levels of parallel computational performance. This system has demonstrated over 60 billion floating-point operations per second (60 GFLOP/sec) of sustained performance for the production computational fluid dynamics (CFD) code OVERFLOW-MLP (13 times that of a 16 CPU C90 system). This system offers even higher performance potential for molecular dynamics simulations.

Recently, the Lomax system was used as the parallelization testbed for the COSMOS ab initio

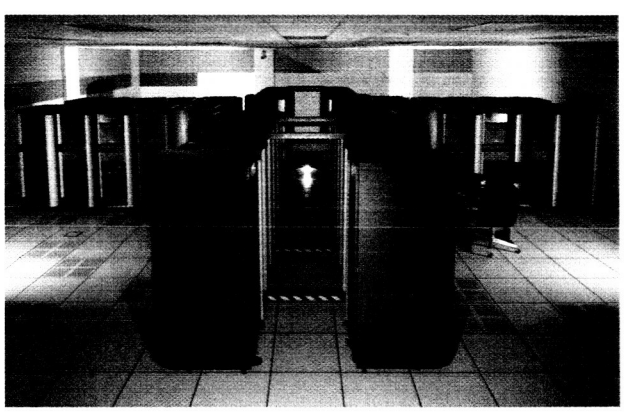

Fig. 1. The Ames 512 CPU SGI Origin 2000 system.

molecular dynamics model used in NASA's astrobiology research effort. The COSMOS code is often used to perform protein-folding simulations. Historically, many important problems involving 20,000-30,000 atoms have not scaled well on "clustered" parallel systems. This lack of performance is due to the small amount of work performed by each CPU relative to the time spent transferring data between CPUs.

The single-system approach of the SGI Origin 2000 architecture, and the large CPU count Lomax system in particular, offers an ideal platform for such computations. The Origin design supports very fast and low latency memory access times from any processor to any memory module. This low latency and high performance are essential for parallel scaling to the hundreds of CPUs necessary to execute problems in a timely manner.

The optimization effort is focused on inserting the highly efficient Ames-developed multi-level parallelism (MLP) approach into COSMOS. At this point the two major time-consuming routines have been converted with highly encouraging results. The first routine computes its zones between all water molecules in the system (WATNLS1). The second (MPFGATHER) gathers the forces for subsequent molecular movement. The results are summarized in Table 1.

Table 1. A comparison of COSMOS and COSMOS-MLP execution times.

COSMOS (32 CPUs)	COSMOS-MLP (343 CPUs)
Module Summary	Module Summary
WATNLS1: 56.66	WATNLS1: 0.94 (60x)
MPFGATHER: 42.13	MPFGATHER: 0.11 (383x)
BARRIER: 0.08	BARRIER: 1.97
Totals: 98.87	**Totals: 2.92 (36x)**

As the table shows, the MLP modifications dramatically improve the code performance on the two most time-dominating routines. The speedup arises from the much higher scaling efficiencies found in the MLP based parallel algorithm, coupled to a greater reuse of encached data. It is this expanded cache reuse that fuels the observed dramatic superlinear speedup over the old code executing at its parallel limit of 32 CPUs.

Current efforts indicate that COSMOS-MLP executions on Lomax will be some of the fastest ever achieved in this field. The results of this research have far-ranging implications in the commercial world, for the advanced numerical techniques developed under this effort are generally applicable to a number of industry standard models used by the university and drug research communities in the United States.

Points of Contact:
J. Taft (COSMOS-MLP)/A. Pohorille (COSMOS)
(650) 604-0704/5759
jtaft@nas.nasa.gov
pohorille@raphael.arc.nasa.gov

Space Technology and CFD Applied to the Development of the DeBakey Heart Assist Device

Cetin Kiris, Dochan Kwak

Approximately 20 million people worldwide suffer annually from heart failure, a quarter of them in America alone. In the United States, only 2,500 donor hearts are available each year. The DeBakey Ventricular Assist Device (VAD) prolongs life until a suitable transplant heart is available, and is used to boost blood flow in patients suffering from hemodynamic deterioration, that is, loss of blood pressure and lowered cardiac output.

The use of computational fluid dynamics (CFD) technology led to major design improvements in the heart assist device, enabling its human implantation. The DeBakey VAD is a miniaturized heart pump designed to increase blood circulation in heart-failure patients awaiting a transplant. A ventricular assist device has to be small and efficient, generating a 5-liter-per-minute blood flow rate against